Anagha P. Krishna

anaghapkrishna2002@gmail.com — Citizenship: US Citizen

Education

Boston University Massachusetts, USA
Masters in Artificial Intelligence *Sep 2025 – Present*

BMS College of Engineering Bengaluru, India
Bachelor of Engineering, Department of Machine Learning *Sep 2020 – Aug 2024*

Work Experience

Healthcare Data Analyst Oct 2024 – Mar 2025
EVERSANA *Bengaluru, India*

- Implemented LLM-based solutions.
- Handled intensive compensation, dashboarding and reporting.
- Executed complex SQL and Python coding for analytics.

Academic Projects

Clinical Notes Assistant — Healthcare RAG *Python, RAG, Qdrant, Ollama, Docker*

- Built a clinical decision-support assistant that lets care teams query patient records and clinical notes in natural language, grounding every answer in the underlying documentation.
- Enforced citation-backed responses that link each answer to the exact source note, and abstain rather than guess when no supporting evidence exists — reducing hallucination risk in a patient-safety-critical setting.
- Deployed a fully self-hosted, on-premise stack (Ollama, Qdrant, Docker) so protected health information (PHI) never leaves the hospital environment.

LLM-Based Intelligent Document Query System *Python, LangChain, FAISS, Transformers, RAG*

- Designed a Retrieval-Augmented Generation (RAG) system to answer natural-language queries over large document collections.

Autonomous ML Pipeline using AI Agents *Python, LangGraph, LLMs, XGBoost, SHAP*

- Designed an autonomous ML pipeline where specialized AI agents handle data exploration, modelling, and evaluation using LangGraph.
- Orchestrated agent interactions via a super-agent that uses LLM-based reasoning to decide the next action in the pipeline.
- Applied SHAP techniques to generate summary-level and instance-level model interpretations.
- Automated end-to-end ML workflow — agents performed feature engineering, feature selection, model training (XGBoost) and exported deployable models in PMML and Pickle formats.

Model Monitoring & Drift Detection *Python, Scikit-learn, Statistical Analysis*

- Simulated real-world data drift scenarios to study their impact on ML model performance.
- Implemented statistical drift detection techniques to monitor feature and prediction distribution shifts.

Cross-Lingual Hallucination Drift in LLMs *Python, GPT-4o, LLM-as-a-Judge, NLP, Multilingual Evaluation*

- Investigated whether hallucination drift in GPT-4o-mini varies by task type across English, Spanish, and Swahili using TruthfulQA (factual recall) and XCOPA (commonsense reasoning).
- Built an LLM-as-a-Judge evaluation pipeline with GPT-4o, computing length-normalized hallucination rates and a task-type drift interaction score (Φ) to test resource-level correlations.
- Validated judge reliability via dual-judge agreement (Cohen's κ) with Claude and span-level F1 against human annotations.

Technical Skills

Programming Languages: Python, C++, C, Java

Databases: SQL, MongoDB

Web Technologies: HTML, CSS

Cloud & MLOps: AWS, Docker, CI/CD

Visualization: Tableau, Excel, Matplotlib, Seaborn

Machine Learning: Large Language Models, Agentic AI, Scikit-learn, TensorFlow, XGBoost, GBM, CNN, RNN, KMeans, Transformers, RAG, Prompt Engineering